




<https://algs4.cs.princeton.edu>

GEOMETRIC APPLICATIONS OF BSTs


- ▶ *1d range search*
- ▶ *line segment intersection*
- ▶ *k-d trees*
- ▶ *context*

Overview

This lecture. Intersections among **geometric objects**.



2d orthogonal range search




line segment intersection

Applications. CAD, games, movies, virtual reality, databases, GIS,

Efficient solutions. **Binary search trees** (and extensions).

Overview





This lecture. Only the tip of the iceberg.



Computer Science 451
Computational Geometry

[Bernard Chazelle](#)

Princeton University
Computer Science
Department





<https://algs4.cs.princeton.edu>

GEOMETRIC APPLICATIONS OF BSTs

- ▶ *1d range search*
- ▶ *line segment intersection*
- ▶ *k-d trees*

1d range search


Extension of ordered symbol table.

- Insert key-value pair.
- Search for key k .
- Delete key k (and associated value).
- **Range search:** find all keys between k_1 and k_2 .
- **Range count:** number of keys between k_1 and k_2 .


Application. Database queries.

Geometric interpretation.

- Keys are point on a **line**.
- Find/count points in a given **1d interval**.



```
insert B      B
insert D      B D
insert A      A B D
insert I      A B D I
insert H      A B D H I
insert F      A B D F H I
insert P      A B D F H I P
search G to K H I
count G to K  2
```



Suppose that the keys are stored in a sorted array of length n .

What is the worst-case running time for **range search** as a function of both m and n ?

number of matching keys

- A. $\Theta(\log m)$
- B. $\Theta(\log n)$
- C. $\Theta(\log n + m)$
- D. $\Theta(m + n)$

1d range search: elementary implementations

Unordered list. Slow insert; slow range search.

Sorted array. Slow insert; fast range search.

order of growth of running time for 1d range search

data structure	insert	range count	range search
unordered list	n	n	n
sorted array	n	$\log n$	$\log n + m$
goal	$\log n$	$\log n$	$\log n + m$

m = number of keys that match


n = number of keys

1d range search: BST implementation

1d range search. Find all keys between k_1 and k_2 .

- Recursively find all keys in left subtree (if any could fall in range).
- Check key in current node.
- Recursively find all keys in right subtree (if any could fall in range).

range search [D...Q]



Proposition. Takes $\Theta(\log n + m)$ time in the worst case.

Pf. Nodes examined = { search path to k_1 } \cup { search path to k_2 } \cup { matches }.

$\Theta(\log n)$

$\Theta(\log n)$

$\Theta(m)$

1d range search: summary of performance

Unordered list. Slow insert; slow range search.

Sorted array. Slow insert; fast range search.

BST. Fast insert; fast range search/count.

order of growth of running time for 1d range search

data structure	insert	range count	range search
unordered list	n	n	n
sorted array	n	$\log n$	$\log n + m$
balanced BST	$\log n$	$\log n$	$\log n + m$

use rank() method
(see precept)

m = number of keys that match
 n = number of keys




<https://algs4.cs.princeton.edu>

GEOMETRIC APPLICATIONS OF BSTs

- ▶ *1d range search*
- ▶ *line segment intersection*
- ▶ *k-d trees*
- ▶ *context*

Orthogonal line segment intersection

Given n horizontal and vertical line segments, find all intersections.



Brute-force $\Theta(n^2)$ algorithm. Check all pairs of line segments for intersection.



Microprocessors and geometry

Early 1970s. microprocessor design became a **geometric** problem.

- Very Large Scale Integration (VLSI).
- Computer-Aided Design (CAD).

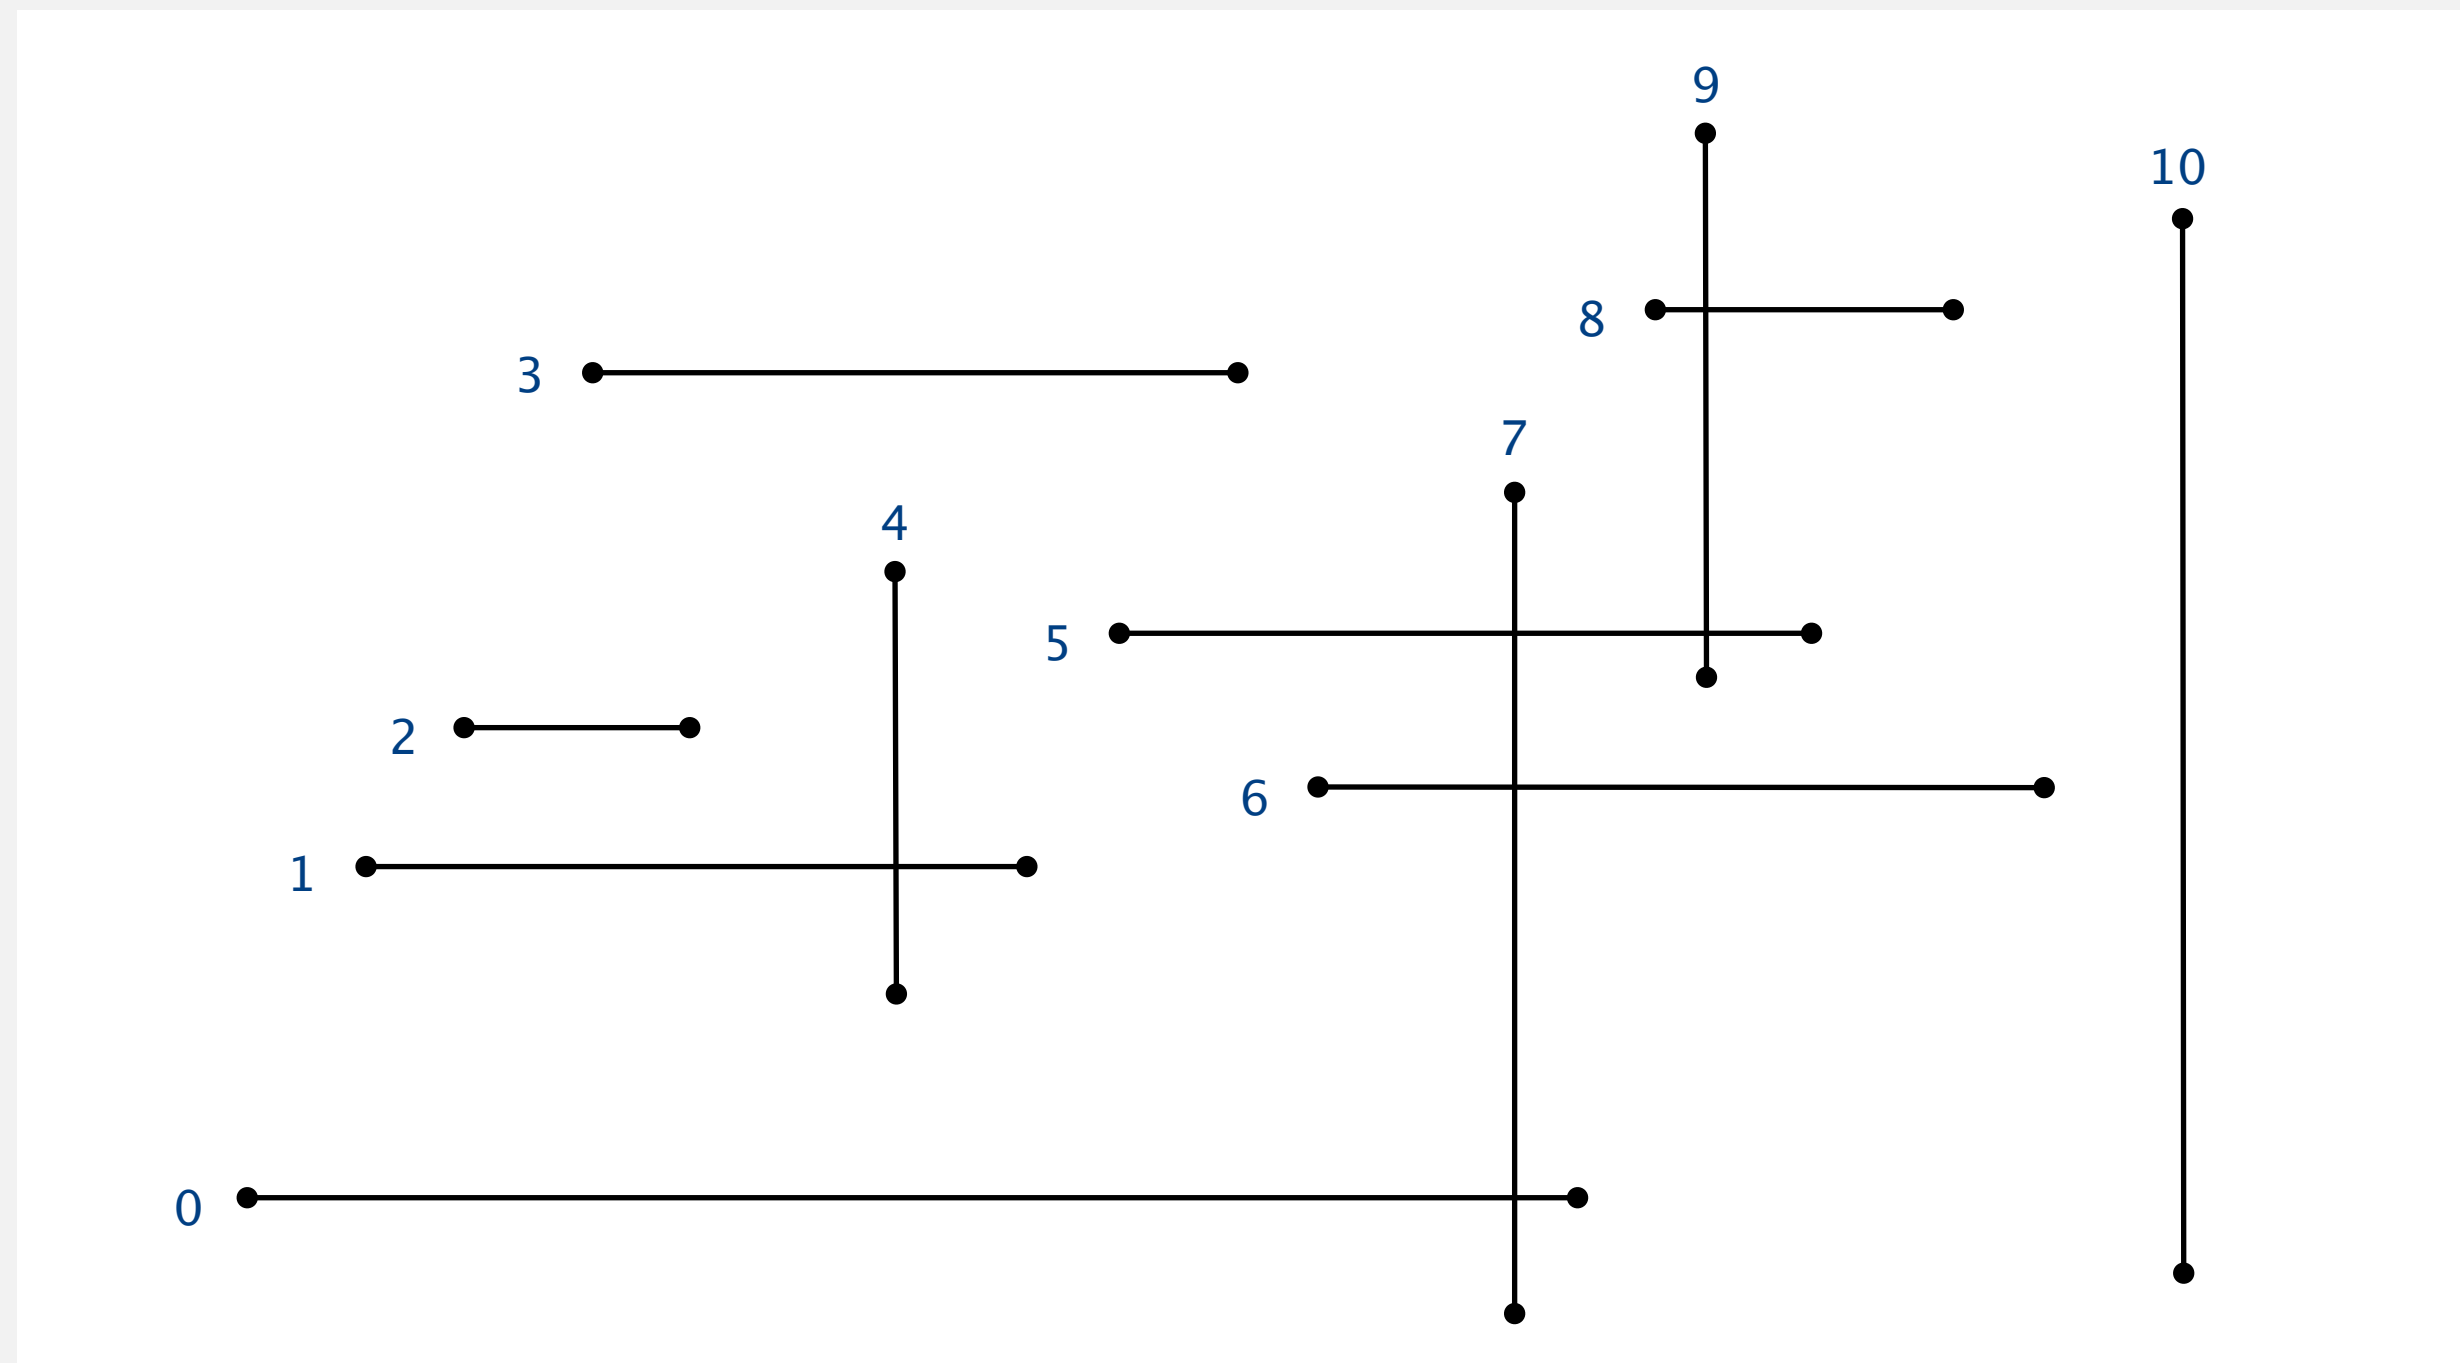
Design-rule checking.

- Certain wires cannot intersect.
- Certain spacing needed between different types of wires.
- Debugging = line segment (or rectangle) intersection.




Orthogonal line segment intersection: sweep-line algorithm

Non-degeneracy assumption. All x - and y -coordinates are distinct.

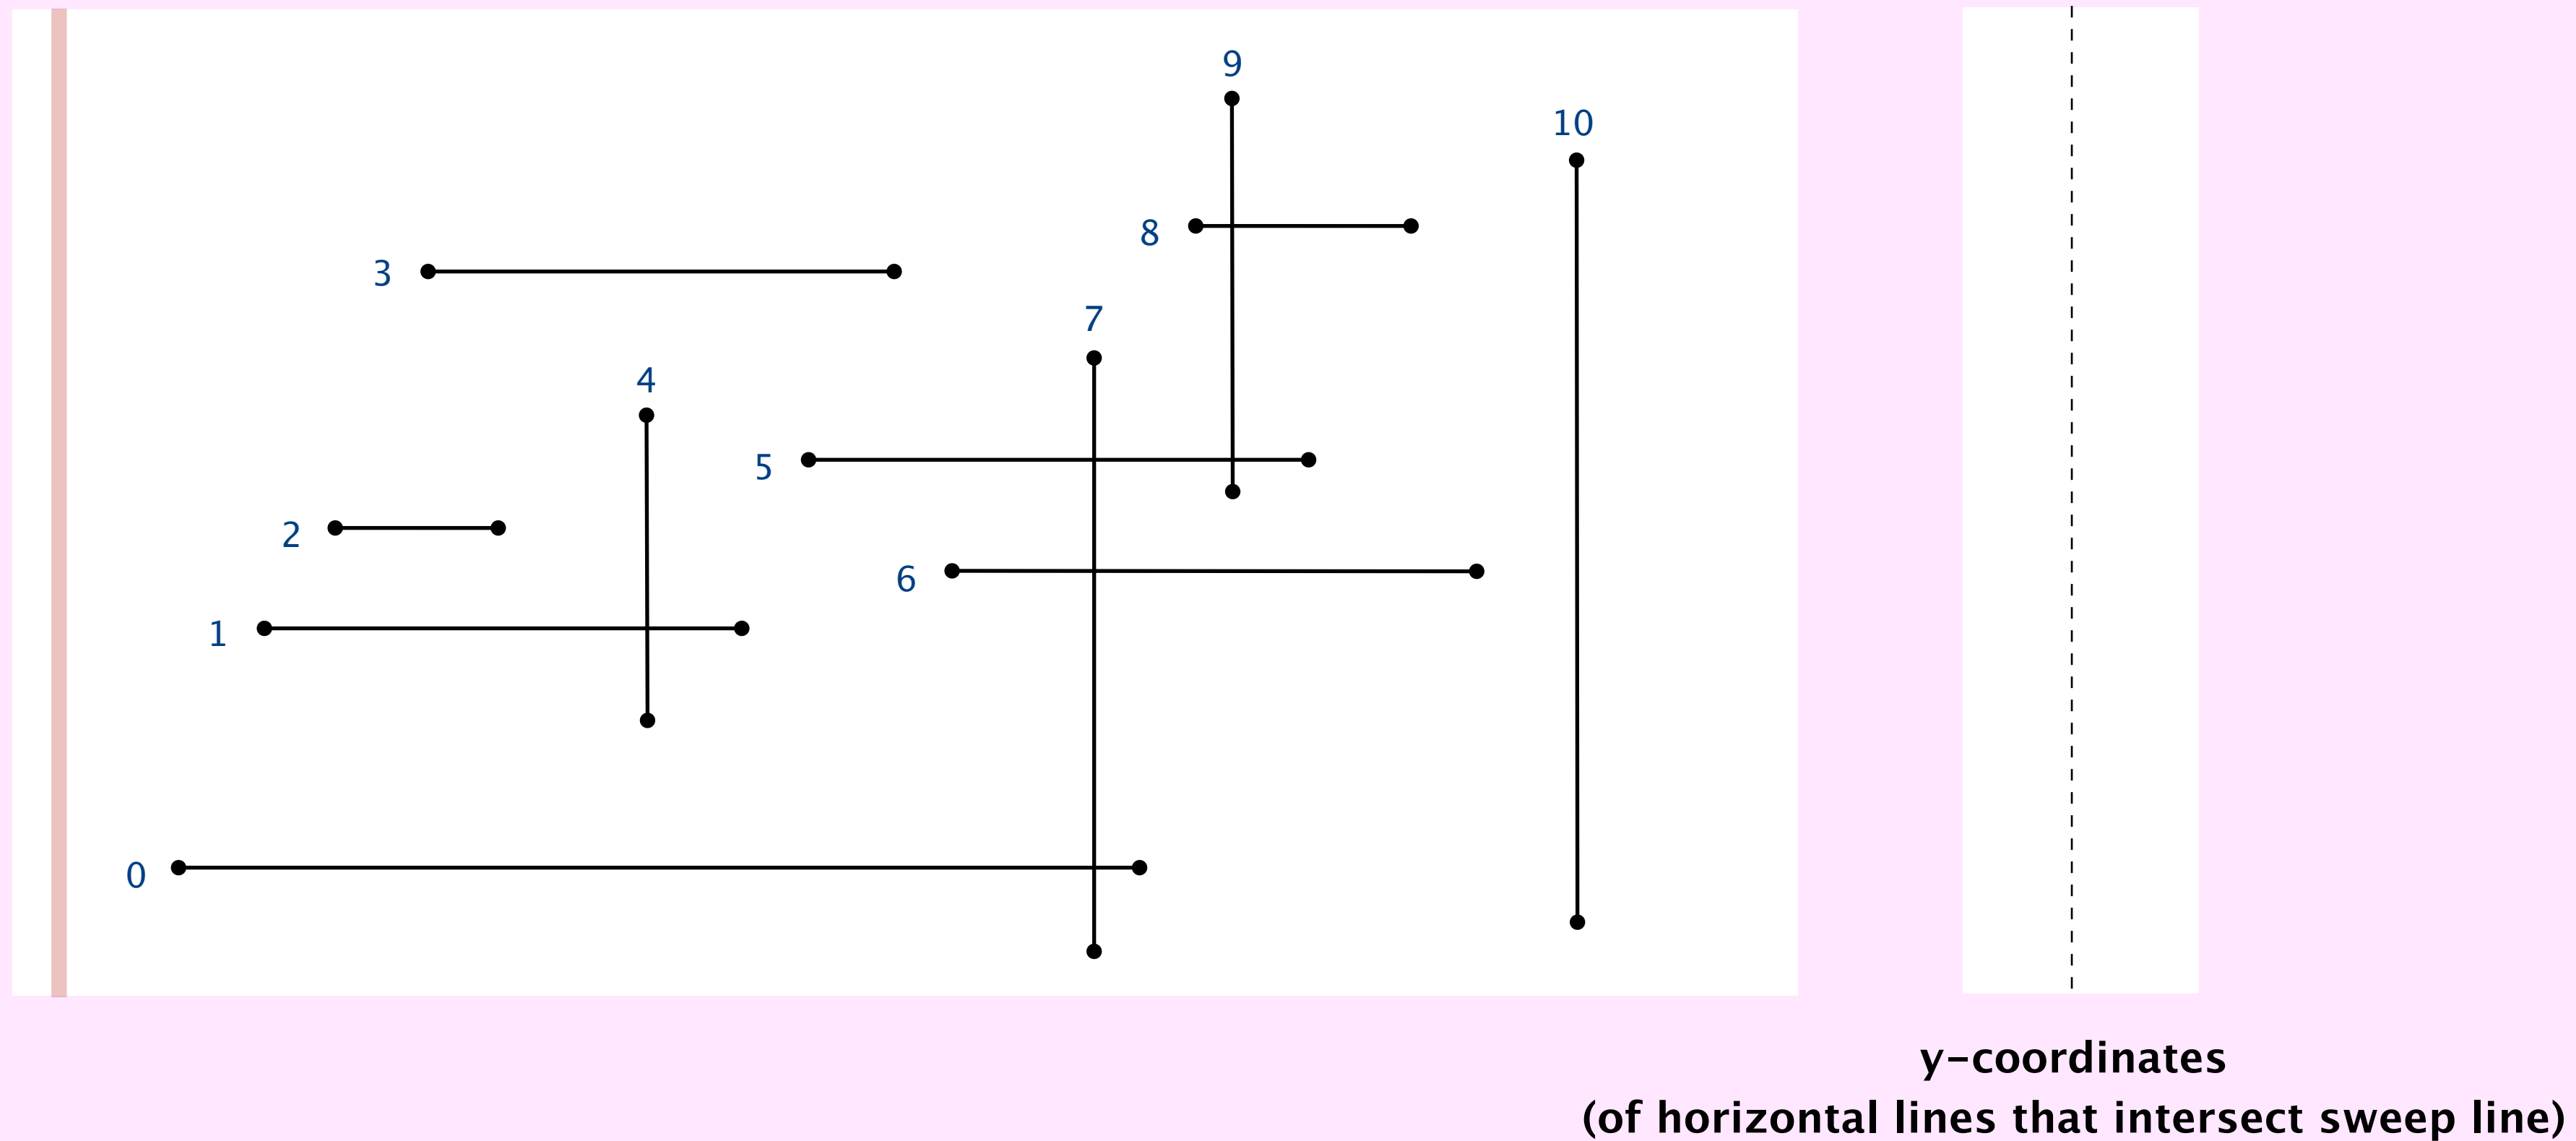


Orthogonal line segment intersection: sweep-line algorithm



Sweep vertical line from left to right. [x -coordinates define events]

- Horizontal segment (left endpoint): insert y -coordinate into BST.
- Horizontal segment (right endpoint): remove y -coordinate from BST.
- Vertical segment: range search for interval of y -endpoints.



Orthogonal line segment intersection: sweep-line analysis

Proposition. The sweep-line algorithm takes $\Theta(n \log n + m)$ time in the worst case to find all m intersections among n horizontal and vertical line segments.

$n \log n$ and m are incomparable
(neither is a lower-order term)

Pf.

- Sort x -coordinates. $[n \log n]$
- Insert y -coordinates into BST. $[n \log n]$
- Delete y -coordinates from BST. $[n \log n]$
- Range searches in BST. $[n \log n + m]$

$$(\log n + m_1) + (\log n + m_2) + (\log n + m_3) + \dots$$

Bottom line. Sweep line reduces 2d orthogonal line segment intersection to 1d range search.

Sweep-line algorithm: context





The **sweep-line algorithm** is a powerful technique in computational geometry.

Geometric intersection.

- General line segment intersection.
- Axis-aligned rectangle intersection.
- ...

More problems.

- Andrew's algorithm for convex hull.
- Fortune's algorithm Voronoi diagram.
- Scanline algorithm for rendering computer graphics.
- ...





<https://algs4.cs.princeton.edu>

GEOMETRIC APPLICATIONS OF BSTs

- ▶ *1d range search*
- ▶ *line segment intersection*
- ▶ ***k-d trees***
- ▶ *context*

Two-dimensional orthogonal range search

Extension of ordered symbol-table to 2d keys.


- Insert a 2d key.
- Search for a 2d key.
- **2d orthogonal range search**: find all keys that lie in a 2d range.
- **2d orthogonal range count**: number of keys that lie in a 2d range.

Applications. Networking, circuit design, databases, ...

Geometric interpretation.

- Keys are point in the **plane**.
- Find/count points in a given **$h-v$ rectangle**.

↑
rectangle is axis-aligned



Space-partitioning trees


Use a **tree** to represent a recursive subdivision of 2d space.

Grid. Divide space uniformly into squares.


Quadtree. Recursively divide space into four quadrants.

2d tree. Recursively divide space into two halfplanes.


BSP tree. Recursively divide space into two regions.




Grid



Quadtree



2d tree





BSP tree


Space-partitioning trees: applications

Applications.


- Ray tracing.
- Flight simulators.
- N-body simulation.
- Collision detection.
- Astronomical databases.
- **Nearest neighbor search.**
- Adaptive mesh generation.
- **2d orthogonal range search.**
- Accelerate rendering in Doom.
- Hidden surface removal and shadow casting.




Grid



Quadtree




2d tree


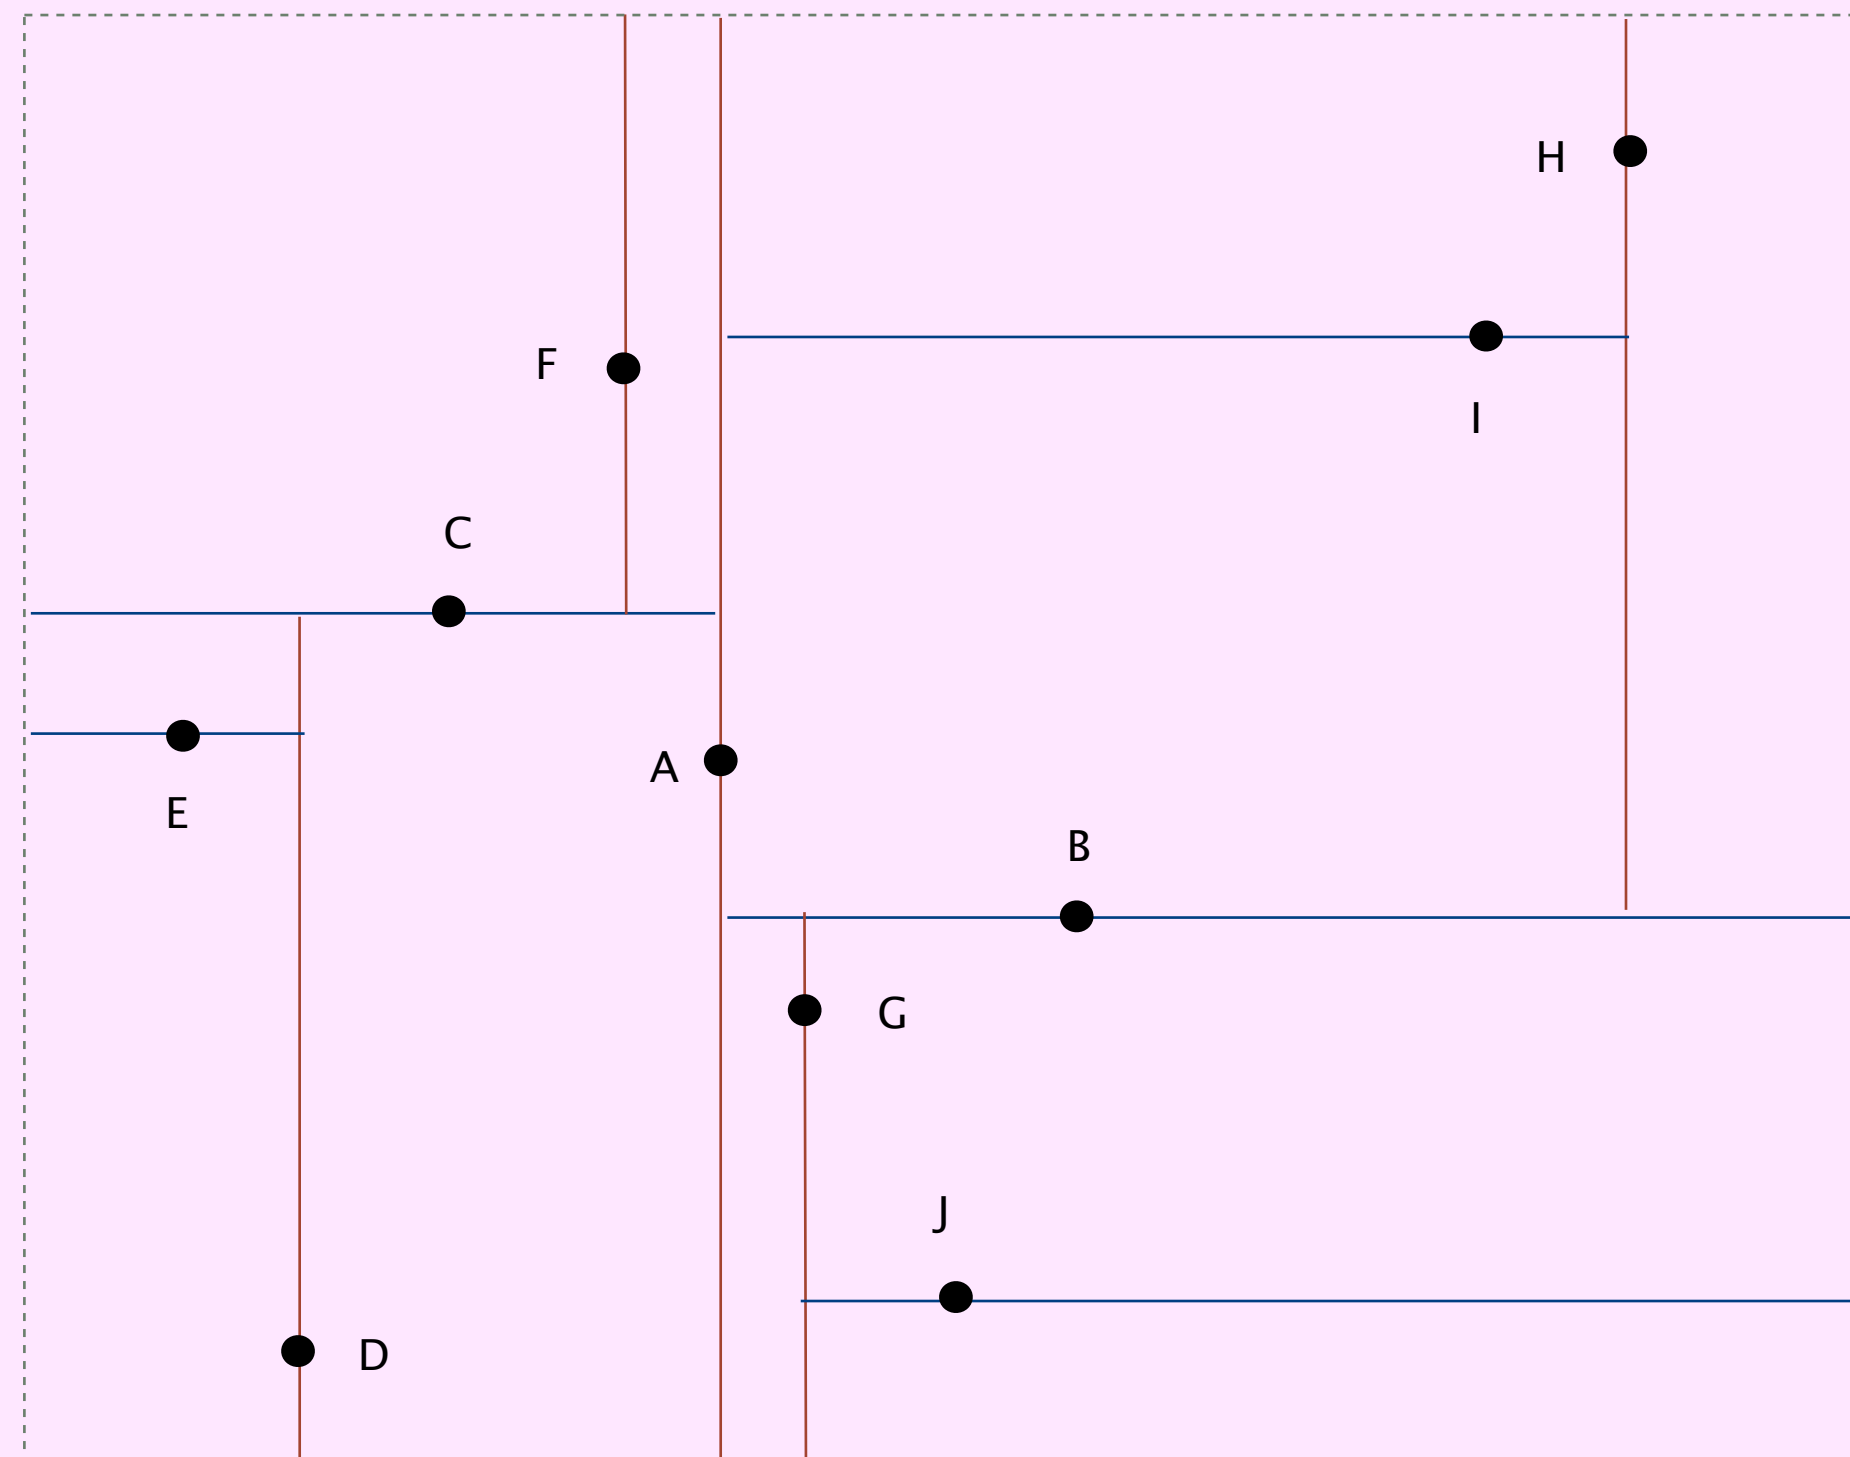


BSP tree

2d tree construction




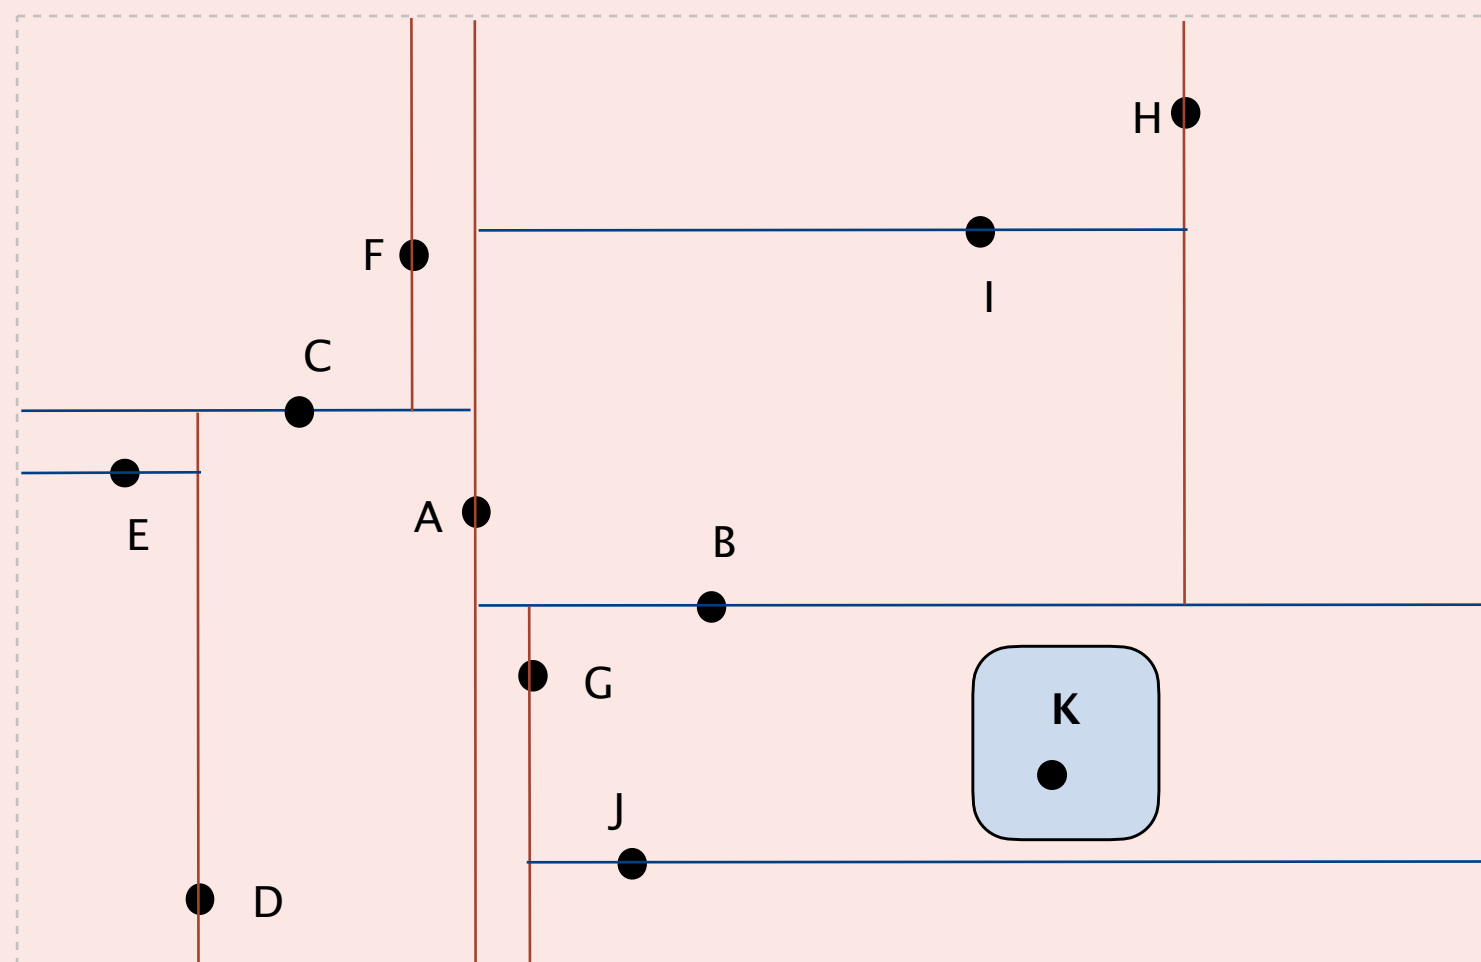
Recursively partition plane into two halfplanes.





Where to insert point K in the 2d tree below?


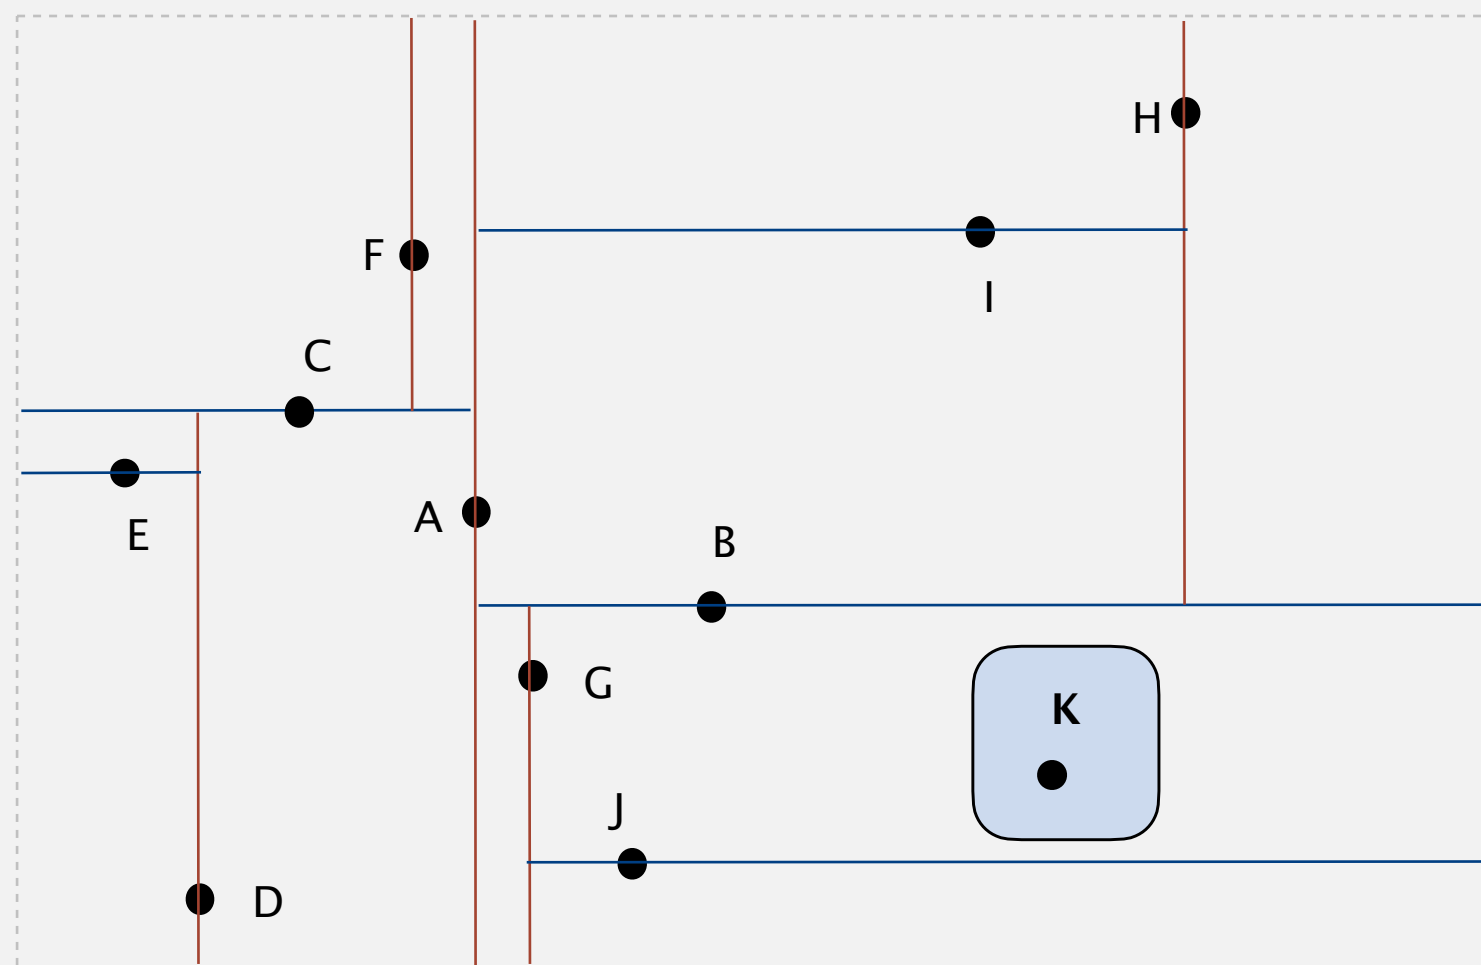


- A. Left child of G.
- B. Left child of J.
- C. Right child of J.
- D. Right child of I.




2d tree implementation

Data structure. BST, but alternate using x - and y -coordinates as key.

- Even levels: compare x -coordinates.
- Odd levels: compare y -coordinates.






2d tree demo: range search



Goal. Find all points in a query rectangle.

- Check if query rectangle contains point in node.
- Recursively search left/bottom and right/top subtrees.
- Optimization: prune subtree if it can't contain a point in rectangle.







Suppose we explore the right/top subtree before the left/bottom subtree in range search. What effect would it have on typical inputs?

- A.** Returns wrong answer.
- B.** Explores more nodes.
- C.** Both A and B.
- D.** Neither A nor B.


Range search in a 2d tree analysis

Typical case. $\Theta(\log n + m)$.


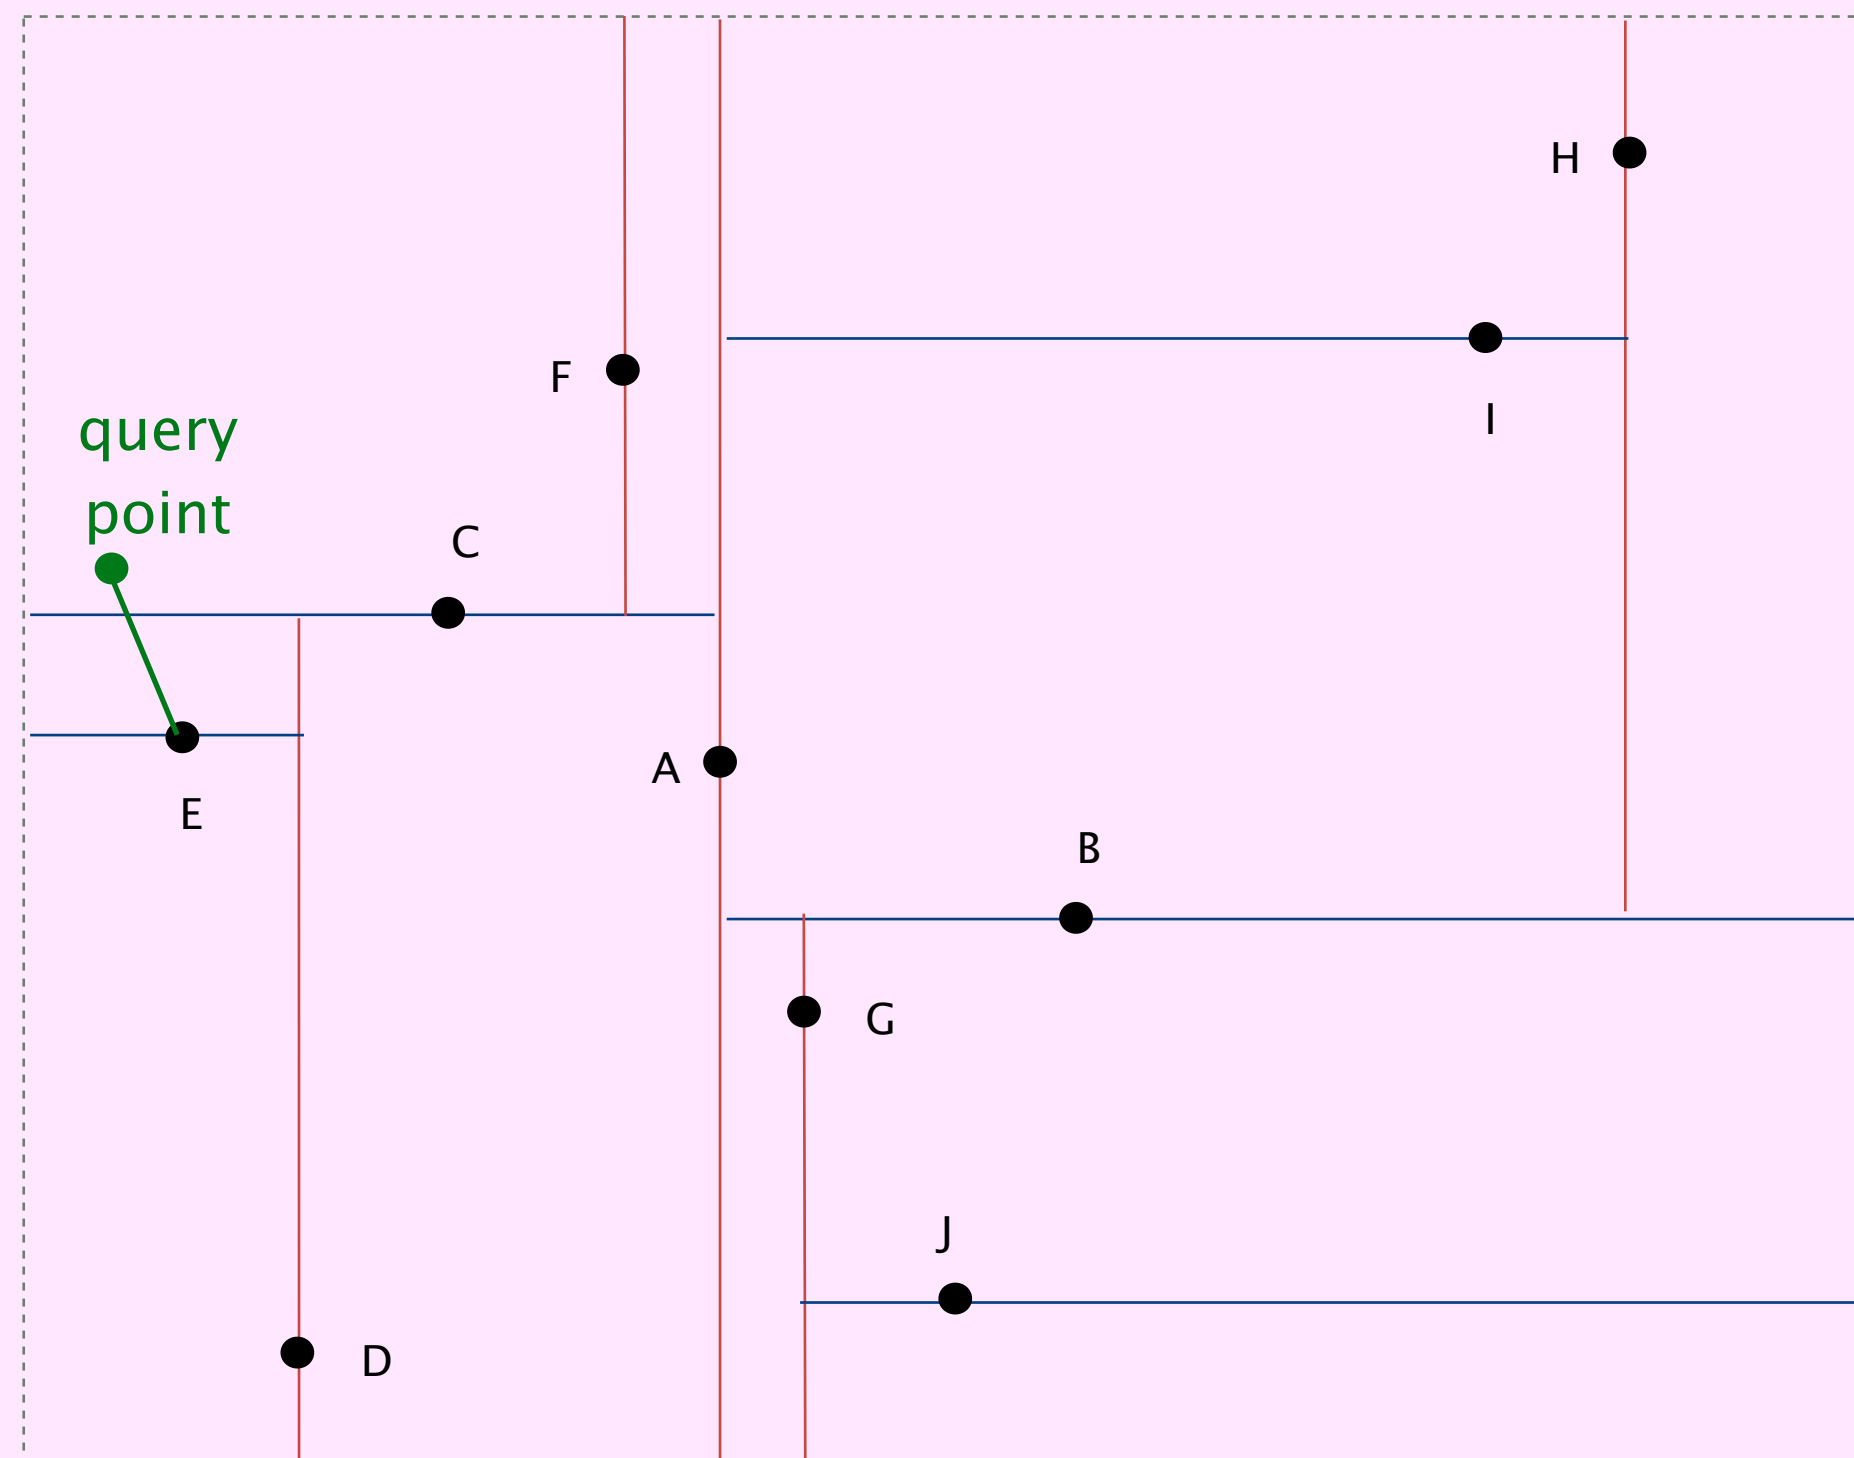
Worst case (assuming tree is balanced). $\Theta(\sqrt{n} + m)$.




2d tree demo: nearest neighbor




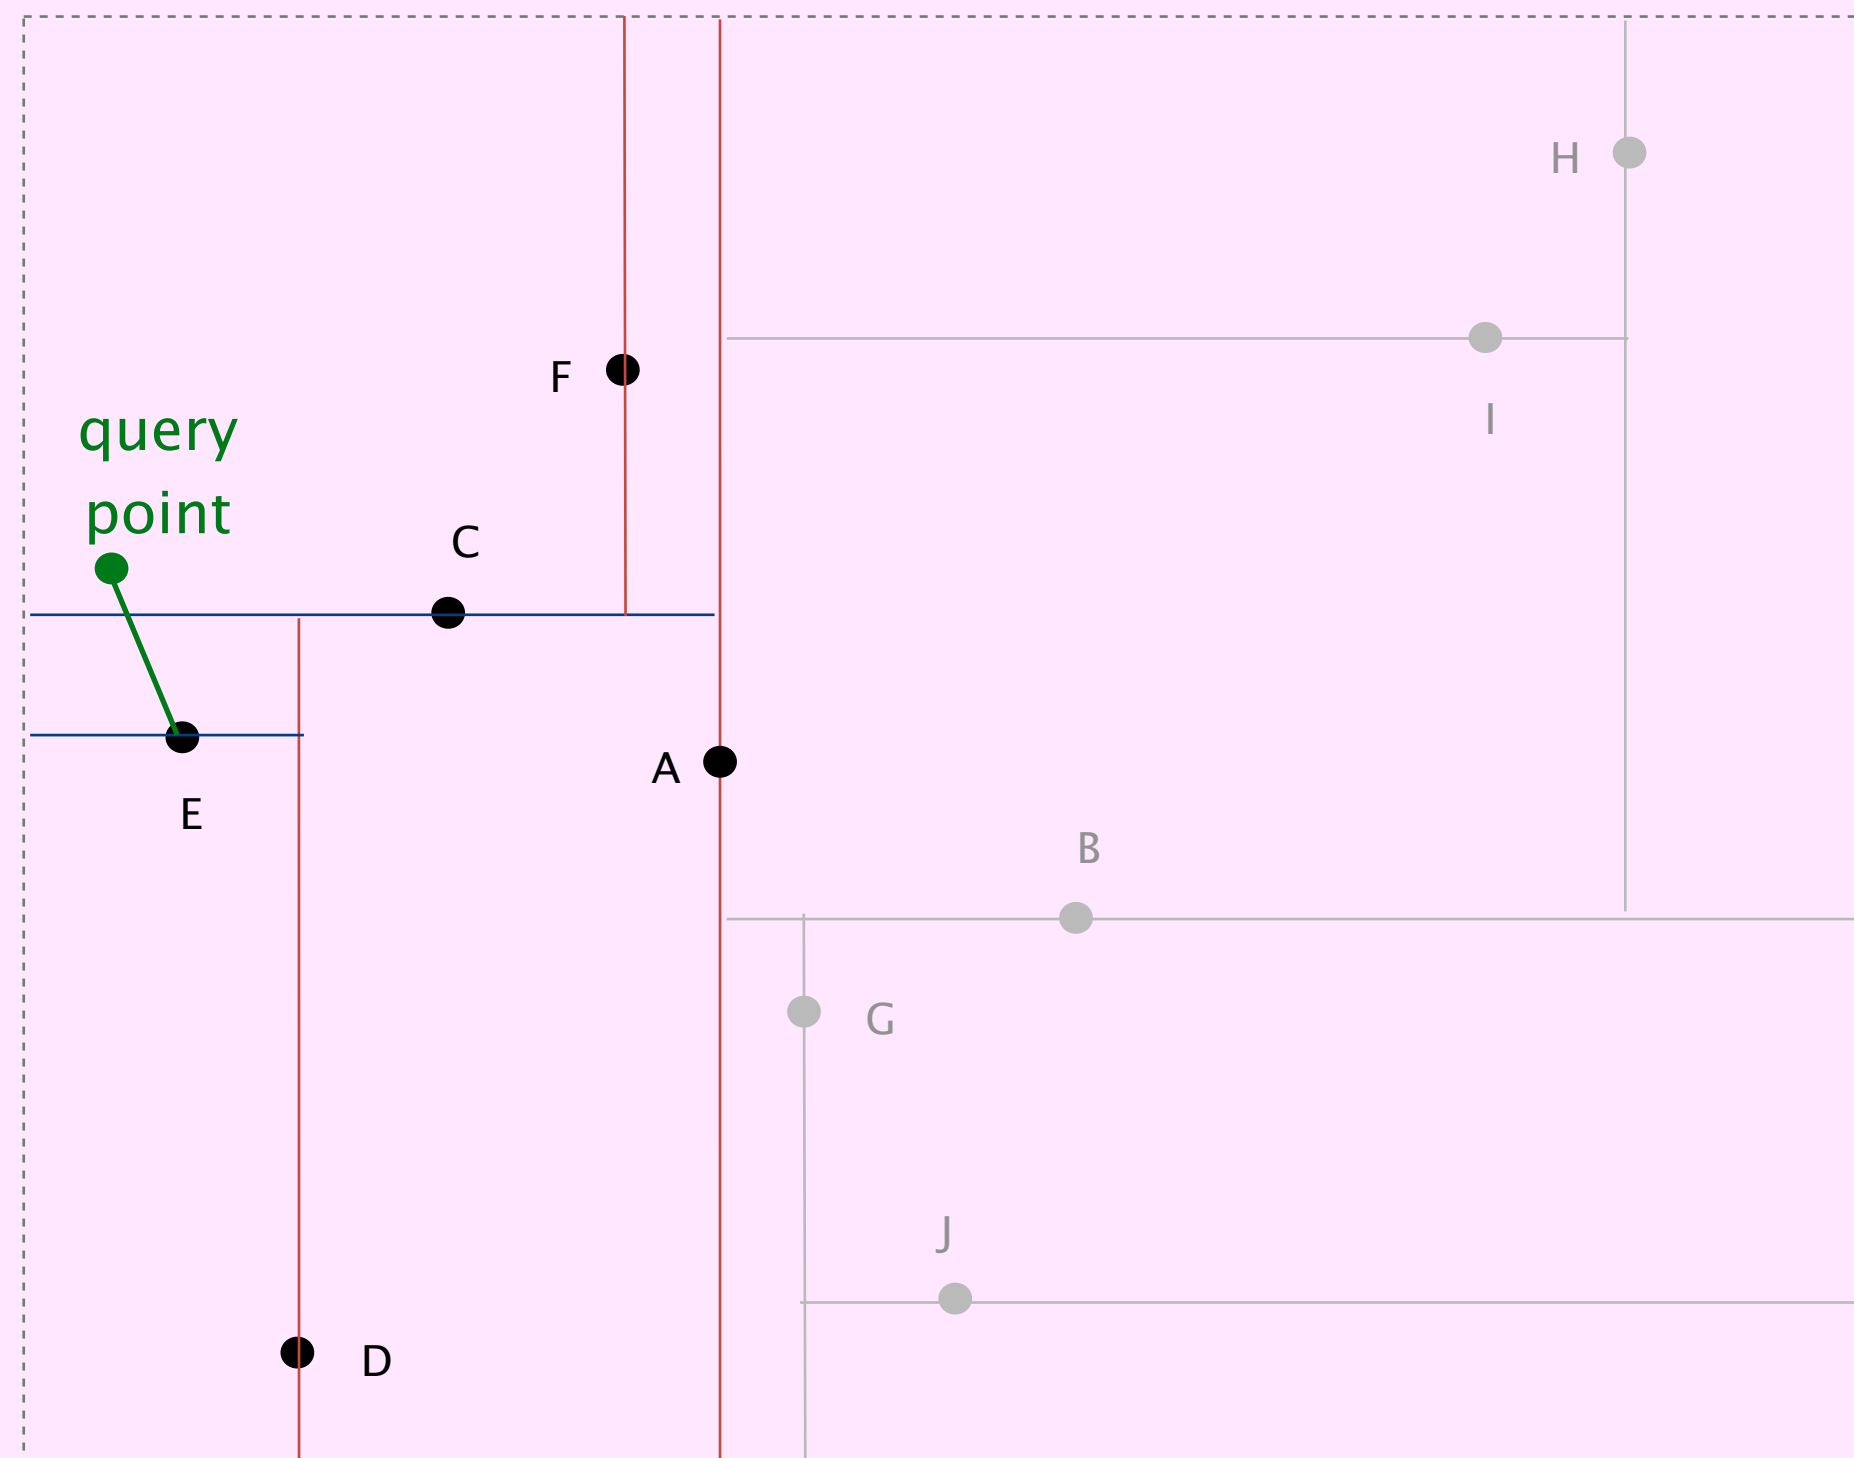
Goal. Find closest point to query point.




2d tree demo: nearest neighbor



- Check distance from point in node to query point.
- Recursively search left/bottom and right/top subtrees.
- Optimization 1: prune subtree if it can't contain a closer point.
- Optimization 2: explore subtree toward the query point first.



nearest neighbor = E



Suppose we always explore the left/bottom subtree before the right/top subtree in nearest-neighbor search. What effect will it have on typical inputs?

- A. Returns wrong answer.
- B. Explores more nodes.
- C. Both A and B.
- D. Neither A nor B.




Which of the following is a worst-case input for nearest-neighbor search?


A.




B.



C.





D.



Nearest neighbor search in a 2d tree analysis

Typical case. $\Theta(\log n)$.

Worst case (even if tree is balanced). $\Theta(n)$.



nearest neighbor = E

2d tree: implementation

Q. How to implement a 2d-tree?

A. Explicit node data type for binary tree.

```
private class KdTreeST<Value>
{
    private Node root;

    private class Node
    {
        private Point2D p;
        private Value val;



        private Node left, right;
        private Node parent;

        private RectHV rect;

        ...
    }
    ...
}
```

very useful for 2D range search
and nearest neighbor search

each node stores a rectangle
corresponding to subdivision of plane



see Assignment 4



<https://algs4.cs.princeton.edu>



GEOMETRIC APPLICATIONS OF BSTs

- ▶ *1d range search*
- ▶ *line segment intersection*
- ▶ *k-d trees*
- ▶ ***context***

K-d tree

K-d tree. Recursively partition k -dimensional space into 2 halfspaces.

Implementation. BST, but cycle through dimensions à la 2d trees.



Efficient, simple data structure for processing k -dimensional data.

- Widely used.
- Adapts well to high-dimensional and clustered data.
- Discovered by an undergrad in an algorithms class!



- Q. Which “natural algorithm” do starlings, migrating geese, starlings, cranes, bait balls of fish, and flashing fireflies use to flock?




<https://www.youtube.com/watch?v=XH-groCeKbE>



Boids. Three simple rules lead to complex emergent flocking behavior:

- **Flock centering (cohesion):** move toward the center of mass of k nearest boids.
- **Direction matching (alignment):** update velocity toward average velocity of k nearest boids.
- **Collision avoidance (separation):** point away from k nearest boids.




N-body simulation

Goal. Simulate the motion of n particles, mutually affected by gravity.

Brute force. For each pair of particles, compute force: $F = \frac{G m_1 m_2}{r^2}$

Running time. Time per step is $\Theta(n^2)$.



https://www.youtube.com/watch?v=ua7YIN4eL_w

Appel's algorithm for n-body simulation

Key idea. Suppose that a particle is in a galaxy far, far away from a cluster of particles.

- Treat the cluster of particles as a single aggregate particle.
- Compute force between particle and **center of mass** of aggregate.



Appel's algorithm for n-body simulation

- Build 3d-tree with n particles as nodes.
- Store center-of-mass of subtree in each node.
- To compute total force acting on a particle, traverse tree, but stop as soon as distance from particle to subdivision is sufficiently large.

SIAM J. SCI. STAT. COMPUT.
Vol. 6, No. 1, January 1985

© 1985 Society for Industrial and Applied Mathematics
008




AN EFFICIENT PROGRAM FOR MANY-BODY SIMULATION*

ANDREW W. APPEL†

Abstract. The simulation of N particles interacting in a gravitational force field is useful in astrophysics, but such simulations become costly for large N . Representing the universe as a tree structure with the particles at the leaves and internal nodes labeled with the centers of mass of their descendants allows several simultaneous attacks on the computation time required by the problem. These approaches range from algorithmic changes (replacing an $O(N^2)$ algorithm with an algorithm whose time-complexity is believed to be $O(N \log N)$) to data structure modifications, code-tuning, and hardware modifications. The changes reduced the running time of a large problem ($N = 10,000$) by a factor of four hundred. This paper describes both the particular program and the methodology underlying such speedups.

Impact. Running time per step is $O(n \log n) \Rightarrow$ enables new research.

Geometric applications of BSTs

problem	example	solution
1d range search		<i>binary search tree</i>
2d orthogonal line segment intersection		<i>sweep line</i> <i>(reduces to 1d range search)</i>
2d range search k-d range search		<i>2d tree</i> <i>k-d tree</i>

© Copyright 2023 Robert Sedgewick and Kevin Wayne



© Copyright 2023 Robert Sedgewick and Kevin Wayne